



Business / Technical Brief

Guide to Artificial Intelligence for Hospitality Executives

Nor1-Oracle Overview of Artificial Intelligence
and Machine Learning

Jason G. Bryant, Vice President Nor1,
Oracle Hospitality

October 2022, Version 1
Copyright © 2022, Oracle and/or its affiliates
Public

VISIT...





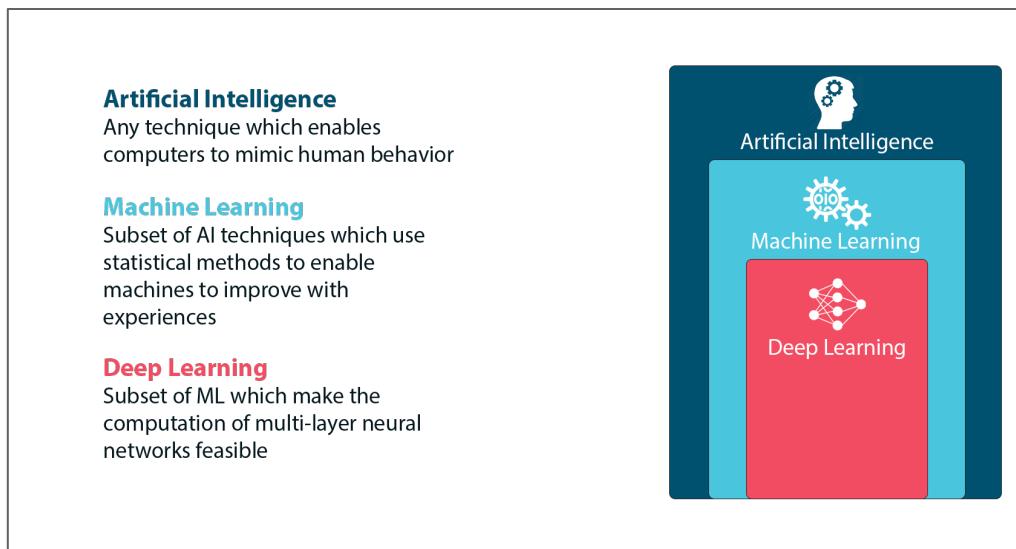
The Basics of Artificial Intelligence and Machine Learning

Recent surveys have shown that fear, ignorance, and misinformation exist regarding artificial intelligence (AI), machine learning (ML), and deep learning (DL). Understanding the relationship between them is the best place to start when exploring how they are used in decision-making technologies. AI can be used not to eliminate jobs, but instead to handle tedious tactical tasks so that staff are empowered to do more strategic work and pay better attention to guests. Jobs will shift, becoming different but no less necessary. This guide breaks down not only the reality of what AI is and does, dispelling myths along the way, but also illuminates the ways it can be used to elevate hotel guest service, create efficiencies, and generate revenue for hotels. As psychiatrist Dr. Karl Menninger once said, “*Fears are educated into us, and can, if we wish, be educated out.*”

The term *artificial intelligence* was coined in the early 1950s when scientists programmed a machine to learn to play checkers. The computer was taught to process data to adjust its outcome, giving it “intelligence.” But because so much data needed to be stored to make a machine intelligent and computers couldn’t yet hold enough information, AI projects languished for decades. Then in the late 1990s, IBM’s Deep Blue defeated world chess champion, Garry Kasparov. Shortly thereafter, computers were trained to recognize human speech as well as human emotions.

According to [Harvard University’s “The History of Artificial Intelligence,”](#) artificial intelligence hasn’t changed all that much over time; what has changed is the capacity of computers to hold and process data, the wide variety of open-source tool kits, and growing academic and industrial expertise in the field of data science. In short, the fundamentals of artificial intelligence are nothing new. What is new is the rapid rate at which we find artificial intelligence making its way into our lives and the general discomfort people have with not knowing what the impact will be.

Much of the anxiety around artificial intelligence, machine learning, and deep learning stems from misunderstandings, and perhaps a bit of fearmongering on the part of the media, which has given consumers the belief that AI will lead to robots taking over society. A whole genre of movies built around this premise—*Blade Runner*, *Total Recall*, *The Matrix*, and many more—whip up anxiety through fictionalized accounts of AI gone wrong. Breaking down the realities and benefits of AI starts with understanding exactly what it is (and is not), separating the fiction from the nonfiction.



The relationship between artificial intelligence, machine learning, and deep learning.

We won't go on at great length about technical aspects of artificial intelligence (*please see the reference section for recommended books that can provide a deeper level of technical explanation*); however, a general understanding and framework can help make AI more accessible and illuminate the practical benefits. Artificial intelligence is the umbrella under which machine learning and deep learning fall. Broadly, AI aims to enable computers to replicate human behaviors. Within AI, machine learning is a type of artificial intelligence that allows algorithms to improve decision-making with experience. The most specialized and sophisticated artificial intelligence, deep learning, imbues machines with neural networks, making complex tasks like self-driving cars a possibility. They give computers the ability to interpret sensory data and are used for programs such as facial recognition and speech recognition.

Basics of Machine Learning (ML)

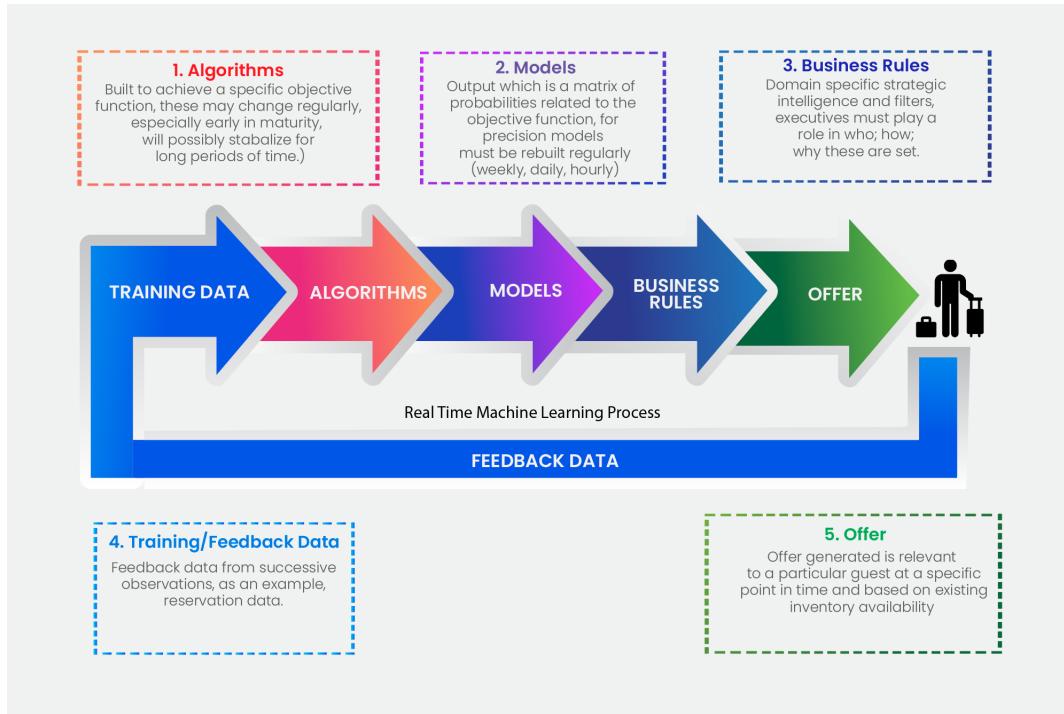
There are many definitions of machine learning. For our purposes with this guide, ML is the scientific study of algorithms and statistical models that computer systems automatically use in real-time to effectively perform a specific objective function (such as optimizing revenue) without using explicit instructions. Instead, ML relies on patterns and inference. All this is performed using a feedback loop so that each successive iteration is used to increase the precision of the models, which drives and improves the objective function.

This is made clear by considering the four primary components of a machine learning system:

1. **Algorithms** - Built to achieve a specific objective function, these may change regularly, especially early in maturity, and will possibly stabilize after a long period of time.
2. **Models** - Output which is a matrix of probabilities related to the objective function. Precision models must rebuild regularly (weekly, daily, hourly).

3. **Business rules** - Domain specific strategic intelligence and filters. Executives must play a role in who, how, why these are set.
4. **Training data** - Feedback data from successive observations, such as reservation data.

ML allows the models to put out smarter information as it learns about consumer behaviors. It does so in a way that humans aren't capable of by creating correlations based on massive amounts of data in real-time, at a speed humans can't replicate.



The machine learning loop for upsell offers.



How the Machine Learning Decision-Making Paradigm is Different

A significant misunderstanding for many executives is how ML is used to make decisions. Typically, humans make decisions based on “causation,” that is, **X** causes **Y** - *If I change this variable, then this outcome will change.*

ML systems on the other hand, make decisions based on “correlation” where a very specific objective function has been defined (such as increased revenue), and each potential option (a presented upsell offer) is valued with a probability of a positive outcome. Over time and with each successive observation (transaction) the accuracy and number of probabilities increase.

To understand machine learning, you must understand the difference
between correlation and causation

Without the ability or resources to rapidly assess a large data source and to calculate in real-time thousands of permutations, “causation” is remarkably accurate. Causation, however, has limitations; it can be difficult to replicate from one human to another, is usually not determined in real-time (a big handicap in today’s digital environment), and works best with significant changes rather than nuances. Many classic experiments have been debunked because they failed to see the big picture of correlation, that is, failing to recognize that any number of factors can contribute to a conclusion. For instance, a study by the National Weight Control Registry found a connection between people who ate breakfast and successful weight loss. But this didn’t mean that eating breakfast caused weight loss. The study failed to control for other factors, such as whether these breakfast eaters worked out or what kind of diet they had. Scientists hadn’t established a causal relationship; they had only identified a correlating factor.

Hotels typically thrive on decision making based on the notion of causal relationships. If I decrease rates \$10 per night during a slow period and bookings increase, then we tend to think that the rate decrease worked. But causal relationships are rarely so simple. Perhaps the rate decrease got the attention of an online channel, and the special rate was shared on social media. In this case, the increased exposure may have improved occupancy rates more than the actual rate change. Humans typically can’t recognize or calculate correlations clearly, but correlations are the foundation of ML. This is the real, undeniable usefulness of ML. When machines do this tactical work, then revenue managers, front desk staff and others are freed to do what they do best—strategize better ways to provide great service to the guest and drive revenue.

ML correlates vast amounts of data points to fine-tune recommendations. This saves hotels from unnecessarily decreasing rates when what was really needed was a different approach—a better understanding of what guests were looking for, where the opportunities were, and what they would be willing to pay and when. Correlating hundreds of data points is something revenue managers and front desk agents cannot do, because humans are only able to make causal associations. But correlating positive and negative historical data (such as the rates that were booked and the rates that weren’t booked) can produce offers that drive revenue rather than undercut rates. Part of the problem with the historical data that most hotels use is that it only accounts for the rates that were paid; it doesn’t track the rates

that weren't booked, and this information (what wasn't booked) is just as important as what was. When positive and negative data are fed into the models, machine learning can give hotels results based on probabilities. This tells the front desk agent which suite upgrade to offer and at what price, based on the probability that this specific guest is likely to book it.

Ignore the scary movies! In reality, ML gives us decisioning capabilities we wouldn't have otherwise. The ability to create probabilities and correlations vastly enhances not only the accuracy of decisions, but also the speed at which they are made. Real-time decisioning on the rates and offers that bring in guests—for without the right offerings, what does a hotel achieve?—allows hotels to refocus their efforts on the hospitality at the industry's core.

Overview of Machine Learning (ML) in the Hospitality Industry

Smarter, Faster and More Profitable Decision Making

The idea that humans are superior decision makers arose with Aristotle, who suggested that humans were the only rational animals. According to recent research, we share decision-making traits with primate relatives, and we're all prone to flawed thinking. Laurie R. Santos and Alexandra G. Rosati of Yale University note that humans “*consistently attend too much to irrelevant information, fall prey to contextual and situational variables, and even rationalize our bad decisions. Moreover, many of these irrational biases operate quickly, effortlessly, and outside of our awareness*” ([NCBI](#)). It's not that we're poor decision makers; we're just occasionally inaccurate and sometimes slow to come to the right conclusion.

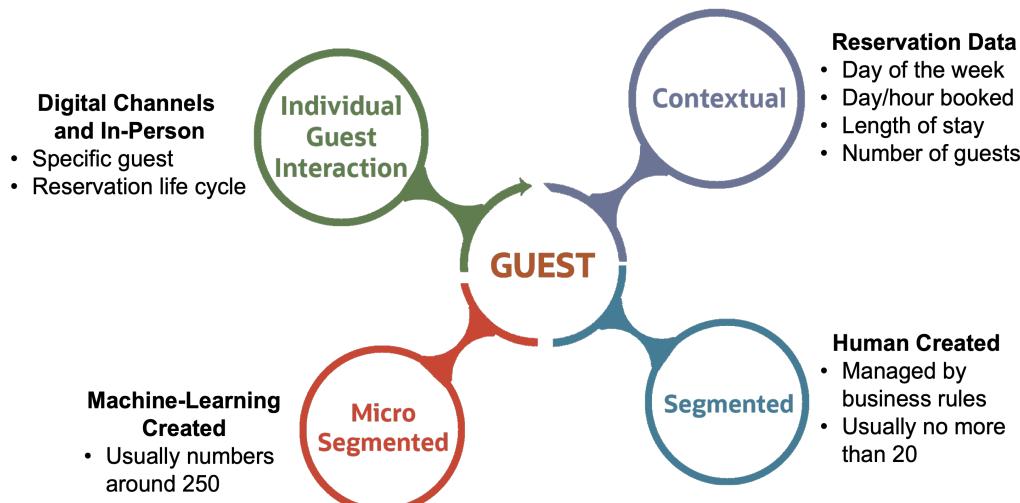
Now, consider what we ask front desk agents to do—respond in a professional but warm and empathetic way to a variety of guests with vastly different wants and needs *and* process their credit card and room assignments as efficiently as possible. At the same time, we want them to understand what the guests might upgrade to without having all the data (or possibly too much data). We like to think that if we give the front desk guidelines and incentives, they'll upsell more effectively but there are too many variables for this to happen.

And what about the many strategic opportunities to reach guests digitally with offers at different points throughout the reservation life cycle? Or the guests who are opting to bypass the front desk with mobile check-in and keyless entry? Given our limited capacity to ascertain what guests want and what they're willing to pay for, as well as the move toward digital-everything, we're due as an industry to reconsider our approach to upselling.

Understanding how to reach the right guest at the right time with the right offer is fundamental to the hotel industry. Automated decisioning came to the travel industry in the 1970s with the introduction of Littlewood's Rule. Littlewood's two-class model introduced revenue management to the airline industry in which discounted fares would be offered if their value exceeded anticipated revenue from the value of full-fare tickets. While ML is far more complex, effective, and efficient than Littlewood's model (which uses static data), it's worth remembering that we've been using algorithms and models for decades to sell to the right guest at the right time for the right price. Now, we're able to process unprecedented amounts of historical data in as little as 70 milliseconds to give the guest the just-right price, the one that the guest will be most willing to pay.

Machine Learning and Personalization

Most AI that today's consumers encounter is driven by an ML approach, and it's usually some form of personalization. Personalization has many definitions and can be achieved in many ways. What passes for personalization at one company may be nothing more than getting the guest's name right, whereas at Nor1 it is the application of a sophisticated data science platform that optimizes in real-time for product and price.



Nor1's Machine Learning Personalization Approach

At Nor1, personalization is achieved when an offer is made real-time, specific to a guest and based on that guest's interests, geography, point in the reservation life cycle and specific purpose of travel. Personalization is an ongoing progression that evolves from making context-based decisions to a micro-segmented perspective of the guest and, ultimately, to the specific guest based on that individual's preferences, clicks, and real-time interaction through the reservation life cycle. At Nor1, we call the consistent, relevant, targeted communications with a guest as "*communication coherency*".

Coherency drives conversion

Few decision systems, even the most sophisticated, can consistently achieve individual guest personalization. The basic constraint is scarcity of data. For instance, this may be the first interaction with this guest (such as if the guest is researching rates online). And even with a known guest, there may not be enough knowledge or data about this specific guest and this specific trip to personalize an offer and the pricing. It's not uncommon for a hotel to have 50% or less repeat business.

At Nor1, our approach to personalization accommodates this scarcity of data with ongoing use of ML integrated into a decisioning platform. Initially, offers and interactions are first driven by categorizing a guest into one or more microsegments. These microsegments change in real-time and are defined by very narrow correlations; they are mathematical representations of the guest at a specific point in time and are like the segmentations used in marketing, like business or leisure, or loyalty program level. The difference of microsegments and typical marketing segments are the number of segments that can be tracked. Usually, we reference 250+ microsegments and an even bigger number of variables used to define the microsegment.

Microsegments are mathematical representations of the guest at a specific point in time and are like the segmentations used in marketing

Once we begin to receive feedback from a guest during the reservation life cycle, we have a greater number of data points to decision on. Personalization at this level is correlating many different things about a guest all at the same time to understand what combination of attributes that guest values, at what price, and at what point the guest will be most likely to upgrade or request services or products. The PRiME Platform and its application of ML surfaces this intelligence during every digital interaction with the guest and then is also provided to front desk agents, empowering them to have a highly relevant and targeted offer for the guest upon check-in.

As an example, when a guest is checking in, the PRiME engine will take into consideration what was already offered to the guest earlier in the life cycle (at booking) and how the guest reacted, for instance did they view, select, or reject the offer, all important regarding what the next offer should be and at what price. In fact, it's the closest we can get to true personalization, using buying signals specific to an individual guest.



A simplified view of how Nor1's PRiME engine works.

A basic characteristic of a real-time machine learning approach like PRiME, where training data constantly updates the models, is that changes in inventory (such as a remodeling or an increased demand during a particular date range) will be automatically recognized and will influence future offers.

PRiME considers where the guest is in the journey, delivering different opportunities at multiple points including booking, pre-arrival, check-in, check-out, and post-stay.

The PRiME decision platform is an example of an *applied AI approach*. Significant effort and investment into PRiME have been made (and are ongoing) to build a tool that specifically focuses on the relationships between guests and the hospitality products they consume throughout the reservation life cycle.

The easiest way to define **applied AI** is to list the five characteristics that creates an **applied AI approach**:

1. Hospitality Industry Specific

When PRiME was first built there were no “off the shelf” AI/ML tools or platforms so the option of buying some generic tool didn’t exist, but now there are many. These generic and off-the shelf systems will never be as accurate or account for industry-specific nuances as a customized system such as PRiME . In fact, we are already seeing how many companies and some of our clients suffer as AI projects are not generating the returns they had hoped – to a large degree because they are taking a generic approach.

2. Produces Results Real-time

Real-time is one of the fundamental features that drives value for any AI system. That is, the ability to look at different signals and data sets, create relevant calculations and then to deliver an intelligent communication or offer to a guest in a split second. On average PRiME creates these calculations and delivers offers to the guest in 70 milliseconds.

3. Contains Continuous, Automated Feedback Loops

A key feature of any ML system is a constant and automatic feedback loop, so the accuracy of the system is self-sustaining.

PRiME gets smarter with every transaction and learns from each guest interaction, including what the guest didn’t select because that’s significant, too. The guest reservation plus their buying profile (responses to offers within this same reservation) combined with a property’s inventory status (to maximize perishable inventory) and historical transactions (hundreds of millions of them) allows ML to serve the ideal targeted offer(s) - the ones most likely to be purchased and generate maximum revenue.

4. Contains the ability to create business rules and other business conditions

Business rules are used to create customized specific strategies to individual organizations (at the corporate, regional, property level). Examples are rules associated with loyalty programs or certain market segments. These must be flexible and easy to modify on a regular basis.

5. Provides Transparent Evaluation of Strategies

The platform must include a tool that provides users (non-data science users) with not only a lens into answering the question of what happened (descriptive) which is what most business intelligence programs do, and with the ability to answer the question “why” did it happen (diagnostic). Users can quickly understand how their strategies are performing and what steps they need to take to improve them. Nor1’s PRiME platform provides this facility with the PRiME Insights tool.



Upselling Across the Guest Journey

By understanding upgradeable room types and attributes, and monetizable products and services, then offering them throughout the guest journey, hotels create more opportunities for upselling. When options are broken down throughout the journey, hotels help guests get more of what they want when the guest is ready to buy. Guests' mindsets about what they want and how much they're willing to spend varies at different times in their journeys. At booking, they are committing to a high spend; many may be reluctant to upgrade at that point. Six weeks later at check-in, they may be more willing to spend to upgrade to a spacious suite or a balcony or book a spa treatment or a special dinner.

Productive upselling across the guest journey relies on pricing and merchandising dictated by a sophisticated algorithm. Showing clearly that the upgrade is being offered at a personalized rate can be an essential component of merchandising the offers to certain guests. When a junior suite is offered at \$69 instead of \$89, the guest can recognize the value immediately and they are more apt to choose the upgrade. Similarly, creating urgency by letting guests know there are only a few of the rooms they want available (e.g., only four rooms still available for these dates) is effective, as it addresses known buyer behavior.

Booking

The airline merchandising models of recent years have shown that travelers are willing to upgrade a certain amount for a better experience. For hotels, that can mean a better room, early check-in, late check-out, cocktails, breakfast buffets, spa, golf, and more. In the past, hotels have presented all the options at booking even those that don't make sense for a guest. For example, consider the way hotels have long shown *every guest every available package* when there's no need for a business traveler, for instance, to have the booking experience muddled by a package designed for families. If properties instead rely on ML to suggest upgrades that make sense for a particular point in the booking cycle, not only are they less likely to abandon their purchase, but they are also far more likely to choose an upgrade offer. ML-powered upselling allows hotels to offer (at booking) just an upgrade to a suite and late check-out if these are most likely to tempt a particular traveler at the time of booking. Other opportunities will be offered later in the journey.

Check-In and On-Property

If a room upgrade wasn't accepted at booking, don't offer it again; instead offer something new based on what you've learned about the guest. Check-in, no matter via front desk or mobile, is also the time to introduce opportunities to expand the experience or make it more convenient, such as adding a breakfast buffet for a business traveler, mini-bar cocktails for a couple, in-room pizza for kids and a bottle of wine for parents during a family vacation.

Reaching out during the trip, when guests are in the very midst of the splurge, is persuasive in a way that trying to promote cocktails or dining packages before the stay can't be.

“Deal-focused guests might get tunnel vision and be so focused on getting the best rate that they miss ancillary offers in the booking path. That’s where eStandby is really valuable... to reach that guest in a pre-arrival email when they might be more receptive... and CheckIn Merchandising is our last chance at the front desk to engage a guest who wasn’t paying attention in the booking path. It’s definitely complementary to our other core booking tools

Dave Van Saun
Director, Ancillary Revenue
Great Wolf Resorts

Post-Stay

Many properties look at post-stay communications as merely a chance to gain feedback. Reviews can't be underestimated, and studies show that they can serve as the basis for rate increases as well. But post-stay emails and texts are equally crucial for engaging guest loyalty. Offers for future stays can be informed by guests' first experiences of the property and what the hotel has learned about their preferences for seasonal travel, room location, and activities.

When hotels work throughout the guest journey to engage the guest, those guests are more likely to choose upgrades and ancillaries. Across Nor1's global portfolio, we see unconstrained demand for upsell offers averaging \$28 per reservation, revenue that many hotels are leaving on the table. While upgrades can and should be offered at booking, hotels must think across the entire journey—both digital and physical—to capitalize on the guest's shifting mindset and use past behavior to influence offerings and merchandising. Relying on millions of points of historical data and the resulting probabilities of booking that only ML can calculate, hotels achieve accurate decisions (safely bypassing the human predisposition to erroneous thinking) that lead directly to increased revenue.



Debunking Common Myths and Identifying Barriers to Success of Machine Learning

Although more hotels are exploring how they could benefit from machine learning, many are holding back on investment due to common misconceptions about ML. In a competitive environment that favors early adopters, hotels that hesitate are likely to miss out on guest satisfaction, streamlined operations, and revenue. Here we'll debunk some of the most common myths around ML and AI for hospitality. Additionally, Nor1 has identified classic barriers that create unnecessary friction to implementation.

Myth #1 – Reservations agents, revenue managers, and front desk staff will become obsolete

Articles and posts appear every day in the press and social media trumpeting that robots will take human jobs. The truth is that we are far from the level of deep learning (DL) that would be required for robots to do everything we do—and that there are things humans can do that machines simply can't. Machines can't strategize. What they do well is to create probabilities. They can enhance our human efforts by doing what the human brain doesn't do well, but the human component—especially in hospitality—is essential. **What machines do well is tactical; what humans do well is strategic.**

Companies like Nor1 are solving for the human component from the start. A core question we ask is *how will ML help hotel team members do more of what they do well*, especially serving the guest. In the long term, ML doesn't mean that humans will be replaced; it means that our jobs will evolve, and they will evolve to suit us better as some of the work we find hard or impossible rolls over to ML.

Myth #2 – We have perfected personalization and don't need ML

This myth speaks specifically to the ways hotels perceive ML and its necessity (or lack thereof). Many properties have been gathering data for years, though the data is almost always siloed and fragmented. Hotels are quick to believe that if an email has the guest's name and some preference data (i.e., she played golf with us once, so let's send her a golf offer), this qualifies as personalization.

But a hotel would only know this preference data during or after the guest stay, and it's not always correct (i.e., maybe she tried golf and hated it). What about the significant number of guests who have never been to your hotel (non-repeat)? How do you personalize for those guests? Personalization requires a different level of understanding of the unknown guest and ML is the optimal tool for this.

ML offers the ability to use vast amounts of data to microsegment and create high conversion probabilities in merchandising and rates for unknown guests; it's the most effective way to make personalized offers to an unknown guest. Offers presented with the ideal pricing for the particular guest are where ML does its most important work for increasing revenue.

Personalization is a bridge between RM and CRM. What to offer and at what price
are equally important for guest acceptance of personalized offers

Myth #3 – Our property has collected data for a long time. We have enough to start right away

Most historical data at hotels only contains positive data. For hotels, positive data includes only the rate that was paid, not the rates that weren't booked. For ML to work, it requires not only the offers and prices that converted but also

the ones that didn't. But even without historical data available at a property, hotels can present ML-based offers because of the vast amount user interaction data, that is how guests responded to offers—hundreds of millions of those transactions—that a company like Nor1 has stored. Why is this important? Because mass quantities of historical data are needed to create the correct probabilities to increase guest acceptance and revenue.

Hotels that implement an ML-based upsell program immediately benefit from these millions of transactions because the program will present that will convert. Over time, a property's own positive and negative data are gathered, which will improve the conversion, generating more revenue.

Myth #4: More data is always better

This is one of the biggest myths of machine learning. It's easy to believe that more data means increased precision and better outcomes. Clearly this is the case at times, but in many cases the reverse can be true.

More data is only better only if it's *relevant* data and not just noise. Care must be taken with changing algorithms to recognize new data points. At Oracle Hospitality we leverage a state-of-the-art machine-learning platform that allows our data scientists the ability explore, experiment, and assess every possible data combination for value (signal vs. noise) on an ongoing basis.

As an example, after working with a client for a few years, the client develops a new loyalty and tiering system. The client then begins to provide Nor1 with this information for each guest reservation. If our objective function is to generate more revenue, does this new data point fortify this pursuit or is it ultimately meaningless.

Myth #5: My business does not need an ML strategy

Every organization should consider the potential positive impact of ML on its strategy and investigate how this technology can be applied to the organization's business problems. Ignoring ML is the same as ignoring the next phase of automation and could place enterprises at a serious competitive disadvantage.

Even if your current ML strategy is "ML is not needed," this should be a conscious decision based on evaluation and consideration. As with every other strategy it should be periodically reconsidered and updated according to the organization's needs, its competitors' status, and evolving technology.



Four Guiding Principles of Artificial Intelligence Systems Design in Hospitality

AI and ML will follow a similar adoption cycle as other powerful, transformative technologies—there will be opposition, reluctance, hesitation, adoption, tremendous success, failures, and possible abuses. Hospitality executives will be able to significantly increase the long-term probability of their success by following five guidelines.

These guidelines should be applied to both internal and external ML projects and/or products.

1. Pledge to Use Data Ethically – An Ongoing Responsibility

The first and most important guideline is to use data ethically. Responsible use of ML includes the ethical, moral, and legal impact of data-science based solutions, and understands there must be ongoing care and consideration employed to generate explainable results, assure algorithmic accountability, and eliminate biases as they are discovered.

2. Domain Specificity and Experience is Paramount to Designing an Effective Solution

Hotels are the site of unique operational nuances caused by the complex relationship between the guest, the inventory, and revenue management. While certain ML techniques and decisioning methodologies are interchangeable between industries, any AI technology employed for hospitality must take this into consideration or the decision system will operate sub-optimally, and possibly be useless.

3. Harness the Power of People and Machine

As we move through this technological cycle, the hospitality industry can and should look toward the cultivation of AI and data skill sets outside of their data science programs. Team members should be trained to achieve a productive collaboration with the AI systems as appropriate to their responsibilities. Not only will this yield the most positive results, but also assures your team evolves at the same pace as the technology. This will create a potent combination, just as training revenue managers to handle yield management technology over the past thirty years has done.

4. Build a Real-time, Sustainable, and Decisionable Data Set

World-class AI systems require tremendously large and varied data sets. Responsibility must be taken to steward this data into usable and accessible formats; most existing data sets and sources are not suited to real-time access. The intent is that the data sets created are democratized across the enterprise, enabling everyone, irrespective of their technical know-how, to work with data comfortably, to feel confident talking about it, and, as a result, make data-informed decisions and build customer experiences powered by data.

AI / ML and its countless uses have opened doors for hospitality organizations to take operations to new, more productive, and truly personalized levels. With principles for use in hand, myths set aside, and hospitality professionals who are ready to embrace the shift and doing more of what they do best, hotels will find their business more efficient, competitive, and profitable and still very much about hospitality.



Glossary

Algorithm: a formula given to a computer in order for it to complete a task (i.e. a set of rules for a computer)

Artificial intelligence: a subset of computer science that deals with computer systems performing tasks with similar, equal, or superior intelligence to that of a human (e.g., decision-making, object classification and detection, speech recognition and translation)

Artificial neural network (ANN): a network modeled after the human brain by creating an artificial neural system via a pattern-recognizing computer algorithm that learns from, interprets, and classifies sensory data

Bayesian networks: also known as **Bayes network, Bayes model, belief network, and decision network**, is a graph-based model representing a set of variables and their dependencies

Big data: large amounts of structured and unstructured data that is too complex to be handled by standard data-processing software

Chatbots: a chat robot that can converse with a human user through text or voice commands. Utilized by e-commerce, education, health, and business industries for ease of communication and to answer user questions

Classification: algorithm technique that allows machines to assign categories to data points

Clustering: algorithm technique that allows machines to group similar data into larger data categories

Cognitive computing: computerized model that mimics human thought processes by data mining, NLP, and pattern recognition

Computer vision: when a machine processes visual input from image files (JPEGs) or camera feeds

Data mining: the process of sorting through large sets of data to identify recurring patterns while establishing problem-solving relationships

Deep learning: a machine learning technique that teaches computers how to learn by rote (i.e. machines mimic learning as a human mind would, by using classification techniques)

Heuristic: a computer science technique designed for quick, optimal, solution-based problem solving

Image recognition: the process of identifying or detecting an object or feature of an object in an image or video

Machine learning (ML): focuses on developing programs that access and use data on their own, leading machines to learn for themselves and improve from learned experiences

Natural language processing (NLP): helps computers process, interpret, and analyze human language and its characteristics by using natural language data

Neural networks: see *artificial neural networks*

Optical Character Recognition (OCR): conversion of images of text (typed, handwritten, or printed), either electronically or mechanically, into machine-encoded text

Pattern recognition: automated recognition of patterns found in data

Reinforcement learning: a machine learning method where the reinforcement algorithm learns by interacting with its environment, and is then penalized or rewarded based off decisions it makes

Robotics: focused on the design and manufacturing of robots that exhibit and/or replicate human intelligence and actions

Robotic process automation (RPA): uses software with AI and ML capabilities to perform repetitive tasks once completed by humans

Structured data: clearly defined data with easily searchable patterns

Supervised learning: a type of machine learning where output datasets teach machines to generate desired outcomes or algorithms (akin to a teacher-student relationship)

Turing Test: a test created by computer scientist Alan Turing (1950) to see if machines could exhibit intelligence equal to or indistinguishable from that of a human

Unstructured data: data without easily searchable patterns (e.g., audio, video, social media content)

Unsupervised learning: a type of machine learning where an algorithm is trained with information that is neither classified nor labeled, thus allowing the algorithm to act without guidance (or supervision)

Reference:

<https://learn.g2.com/artificial-intelligence-terms>



Further Reading

We find the following books extremely helpful to understand the world of data science in a coherent, easy-to-understand fashion. These books are highly recommended for anyone looking to further their knowledge and expertise on this topic.

Big Data: A Revolution That Will Transform How We Live, Work and Think

By Viktor Mayer-Schönberger and Kenneth Cukier

© 2013 An Eamon Dolan Book Mariner Books

Human + Machine: Reimagining Work in the Age of AI

By Paul R. Daugherty and H. James Wilson

© 2018 Harvard Business Review Press

The AI Advantage: How to Put the Artificial Intelligence Revolution to Work

By Thomas H. Davenport

© 2018 Management on the Cutting Edge, MIT Sloan Management Review

Data Science (MIT Press Essential Knowledge series)

By John D. Kelleher and Brendan Tierney

© 2018 MIT Press Essential Knowledge Series

The Hundred-Page Machine Learning Book

By Andriy Burkov

© 2019 Andriy Burkov

Machine Learning for Absolute Beginners

By Oliver Theobald

© 2022 Oliver Theobald

Data Science Strategy for Dummies

By Ulrika Jägare

© 2019 John Wiley and Sons

Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2022, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail REVREC_US@oracle.com.
